## Intel's 22-nm Tri-gate Transistors Exposed

*Contributed by Dick James*

Last week Intel had their Q1 conference call for financial analysts, and revealed that the 22-nm Ivy Bridge parts would make up 25% of their shipment volume in the second quarter of this year. That means that a good quantity will already have shipped, and we managed to track some down in Hong Kong a few weeks ago. Of course we got in touch ASAP and the parts duly arrived, and they *were* the real thing.

Fig. 1 Intel Xeon E3-1230V2 Server CPU

Fig.2 Intel Xeon E3-1230V2 Die

We obtained samples of Xeon E3-1230 v2 CPUs, which are four-core, 3.3 GHz, 64-bit parts intended for the server market. A quick cross-section reveals that Intel have stayed with the nine metal layers used in the last two generations:

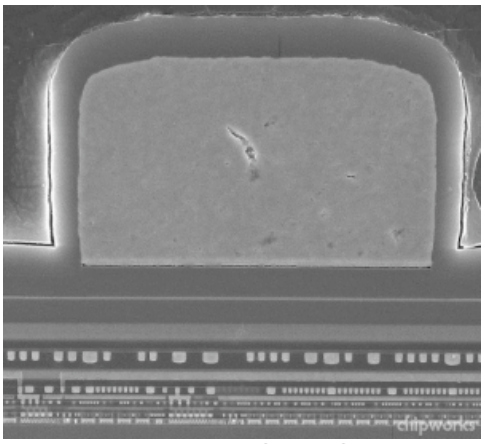

Fig. 3 Intel Xeon E3-130V2 General Structure

A closer TEM image (Fig. 4) shows the lower metal stack and a pair of multi-fin NMOS and PMOS transistors. This section is parallel to the gate, across the fins, and we can see the diamond-shaped epi-SiGe that has been formed on the fins of the PMOS transistor.

We have to digress here a little to explain what we're looking at. A typical TEM sample is 80 – 100 nm thick, to be thin enough to be transparent to the electron beam and at the same time have enough physical rigidity so that it does not bend or fall apart.

Here we are trying to image structures in a die with a gate length of less than 30 nm; so if we make a sample parallel to the gate, and if the sample is aligned perfectly along the centre of the gate, then it will contain the gate plus at least part of the source/drain (S/D) silicon and contacts on either side.

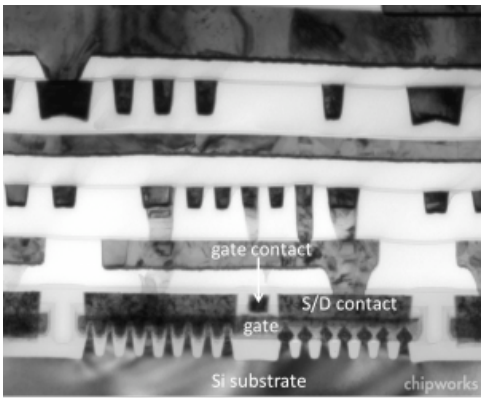


Fig. 4 TEM Image of Lower Metals and NMOS and PMOS (right) Transistors

That is what we see above – I have labeled the gate and contact stripes, and we have PMOS on the right and NMOS on the left. The tungsten-filled contacts obscure parts of the gate, but we can clearly see that the PMOS S/D fins have epitaxial growth on them, and the fins have an unexpected slope – a little different from Intel's tri-gate schematic shown last year – see Fig.5.

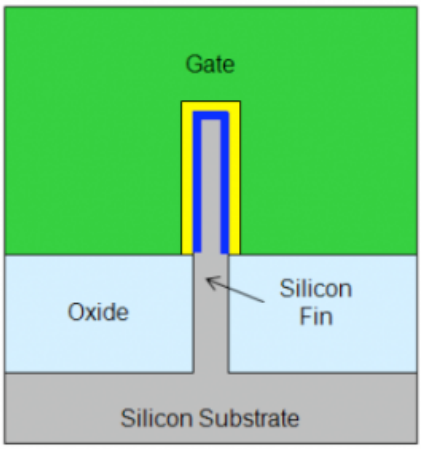


Fig. 5 Intel Schematic of Tri-Gate Transistor

If we zoom in a bit further into the PMOS gate (Fig. 6), we can see how the gate wraps over the fin, and the rounded top of the fin. The thin dark line adjacent to the fin is the high-k layer and just above that is a mottled TiN layer that is likely the PMOS work-function material, as in the 32-nm and 45-nm parts.

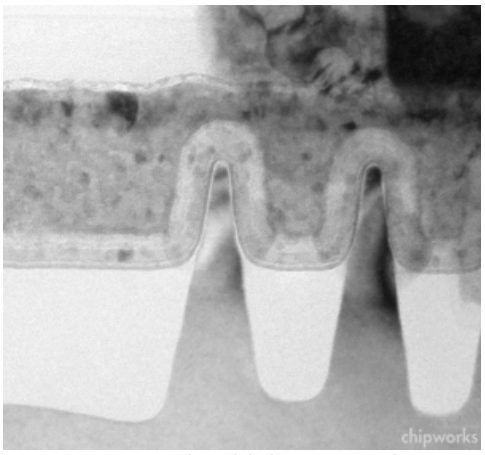


Fig. 6 TEM Image of PMOS Gate and Fin Structure

Fig. 7 shows a section of an NMOS transistor. There is a 'ghost' of the contact behind the gate, but the gate structure itself looks similar to the PMOS, with the exception of the work-function material just above the high-k layer.

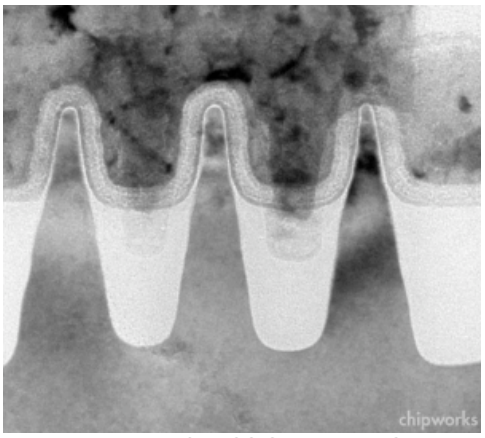

Fig. 7 TEM Image of NMOS Gate and Fin Structure

Fig. 8 gives me an opportunity to show off our new TEM – we have recently purchased an FEI Technai machine, which upgrades our capability considerably. Here we have a lattice image of a fin in an NMOS transistor; the diamond-like layout of the pattern of dots in the fin is created by the columns of atoms in the silicon crystal lattice. This tells us that the sample is oriented in the <110> direction, which given that silicon has a face-centred cubic structure in which equivalent planes are at right angles, means that the channel direction is also <110>.

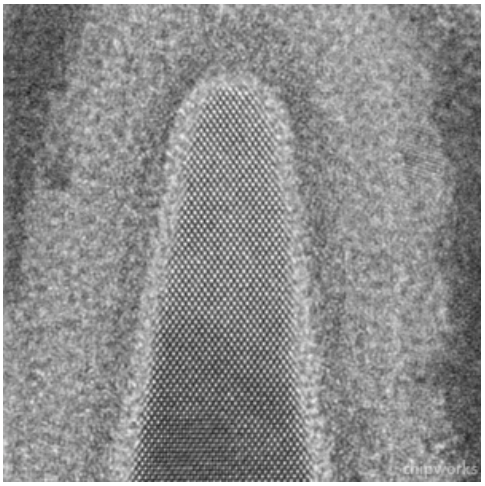

Fig. 8 TEM Lattice Image of NMOS Fin Structure

To fully understand what we're looking at, of course, we need to see what's happening in the orthogonal direction, along the fin and cross-sectioning the gate – as in Fig. 9. This shows an array of PMOS transistors over a single fin, four functional gates and two dummy gates at the ends of the fin. Again the TEM sample is thick compared with the feature size, so we are seeing the gate on the side(s) of the fin, not just the top. The fin ends have the same taper as in Figs 6 and 7.

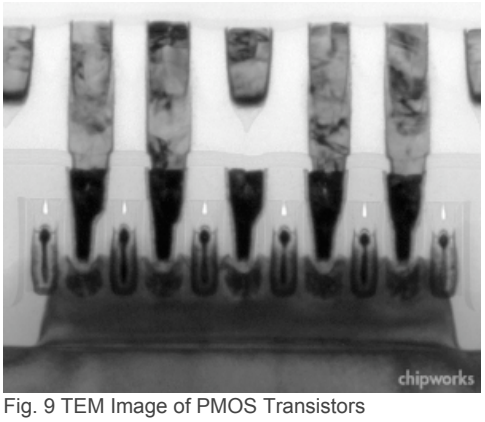

Fig. 9 TEM Image of PMOS Transistors

As announced by Intel, there is embedded SiGe in the source/drains, although not etched to the <111> planes as in the 32- and 45-nm product. It also looks as though the tops of the gates have been etched back and back-filled with dielectric, and the contacts are self-aligned as in memory chips.

Zooming in on the PMOS transistor in Fig.10, the image is a bit fuzzy, but the SiGe is clearly in a rounded cavity with no facets on the top, though there are facets on the sides of the fin (see fig. 4).

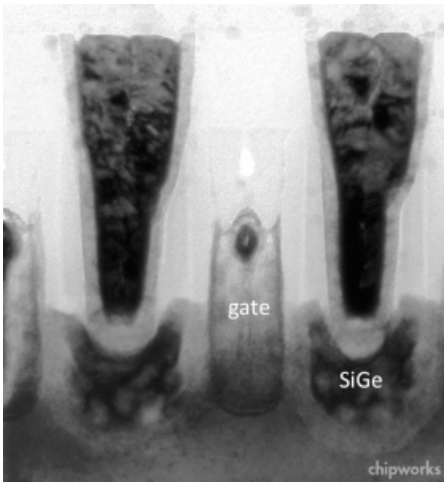

Fig. 10 TEM Image of PMOS Transistor

Looking at the NMOS equivalent (Figs. 11 and 12), we see a similar structure – there seems to be an epitaxial interface, and the silicide(?) seems to protrude slightly above the fin.

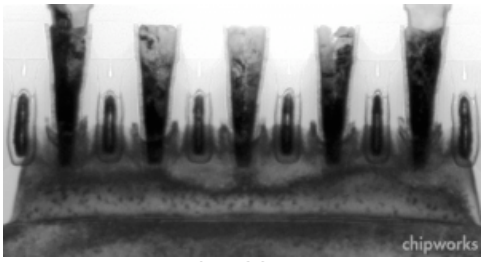

Fig. 11 TEM Image of NMOS Transistors

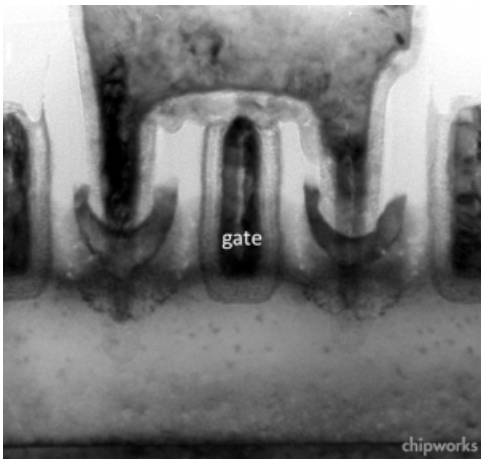

Fig. 12 TEM Image of NMOS Transistors

It is hard to say much about the gates here, either NMOS or PMOS, because of the sample thickness problem; we are viewing a slice that includes the gate on both sides of the fin and the fin itself. Fortunately we have images of gate metal over STI and they are less confusing.

Figure 13 is a composite image of NMOS and PMOS gates so that the differences are more noticeable. The dark line surrounding the gate structures is the Hf-based high-k, and within that are the two work-function materials, likely TiN for PMOS and TiAlN for NMOS. (The columnar structure of the PMOS TiN is visible in the right half of the image.)

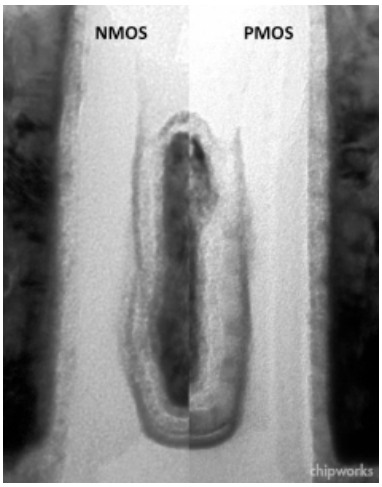

Fig. 13 Composite TEM Image of NMOS/PMOS Gates

The fill has been changed from TiAl in the earlier parts to tungsten. It is more prominent in the NMOS gates than the PMOS, because the PMOS structure includes both work-function metals, whereas the TiN has been etched out of the NMOS gates. At the 45-nm node Intel used tensile tungsten in the contacts to apply channel stress – have they transposed this to the gates in the 22-nm process?

Just to finish up, so that this is still a blog, not a paper (I don't want to go on too long) – fig. 14 shows a sample delayered to expose the transistors, and imaged on a tilt angle. Both the gates and the fins show up nicely, and we can actually see tiny spikes of SiGe in the PMOS source/drains. The small pillars in between the fins in the NMOS areas are residual bits of contact metal. I think it's a cool image!

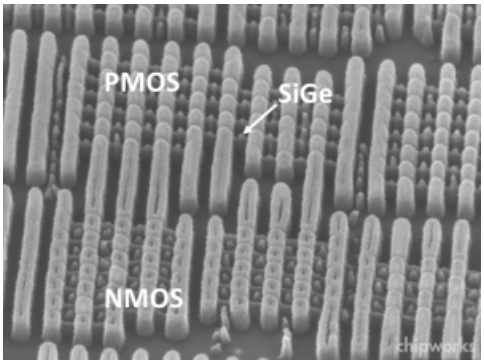

Fig. 14 Tilt SEM Image of NMOS/PMOS Transistors

We are just getting into the full scope of the analysis, so likely more to come in the next few weeks! Below is a link to the set of reports that we are doing on this ground-breaking part, with many more details than I will ever get a chance to write about:

Intel 22-nm Ivy Bridge Xeon E3-1230v2 Microprocessor Structural Analysis, Transistor Characterization, and Package Analysis Reports

Share/Bookmark

This entry was posted on Monday, April 23rd, 2012 at 12:36 pm *by Dick James* and is filed under Advanced CMOS Technology, CPUs. Both comments and pings are currently closed.

## 17 Responses to "Intel's 22-nm Tri-gate Transistors Exposed"

*Wizard_of_Si* says:
April 27, 2012 at 1:44 pm
Great report! Thank you. The TEMs show the complexity of the technology and it is clearer why Intel had to spend north of $15B in capex and process development.

Now the Ivy Bridge power is 140.9W a 6 Watts improvement on Sandy Bridge 146.9W at same frequency. The improvement seems negligible compared to the investment.

http://www.tomshardware.com/reviews/ivy-bridge-benchmark-core-i7-3770k,3181-23.html

The negligible power improvement is perhaps due to the fins' variation? On figure 7, the fins are tapered and not

identical. All fins should be identical to guarantee a minimal variation. How can the manufacturing of the billions of tapered fins be controlled? Looks like 22nm might not provide Intel the power advantage to compete in the mobile SOC market.

- *Hardcall* says:
  May 18, 2012 at 5:43 pm
  I think you don't understand how the economics of scale work for Intel. Intel wants lower power usage mainly for mobile, as barbie points out.

  More importantly, 22 nm chips allow for dramatic shrinkage of die area which saves Intel a lot of money. Since they own their fabs, it's not that it saves them money but that they can make more chips for the same amount of money. This allows them to serve the existing market, meet new demand, and cut costs all at the same time.

  To put this into perspective, Intel sells close to 1 million chips per day. (82 million chips sold in Q4 2011, a 90-day period)

*barbie* says:
May 18, 2012 at 9:46 am
@Wizard_of_Si: who exactly cares a lot about power efficiency of desktops? I'm surprised toms would go measure that really… They even mention that we should wait for the Ultrabook launches to get a better picture of why the evolution makes sense.

- *Dr chip designer* says:
  May 19, 2012 at 10:19 am
  Desktop market is less important but it still represent a very solid data point. As Tom points out laptops are harder to benchmark since so many things change like moving to ssd and ddr3 in newer platform.

  I am ex-process engineer now looking at trigate/finfet from a design perspective. So my interest is does moving to finfet make sense for cost and power. I have real doubts. My guess is nearly all ideal advantages at single transistor level will be lost at product level do to variation in a tri-gate. I think that is what toms data is showing.

*Mary* says:
May 19, 2012 at 12:58 pm
Mr James,

Thanks you for sharing this interesting info.

From RF/analog perspective, I also don't like all the variation in figs. 6 and 7: both fin width and shape and gate oxide becomes thicker at bottom of every single fin and "oxide thickening" is highly variable.

It will be very interesting to watch what chips Intel moves to this structure with all this extra transistor variation.

I would guess Intel's Infineon RF/3G/LTE will continue to use foundry. Likely even move to foundry 28nm next. I would not want to design RF or analog on this.

*jb* says:
May 21, 2012 at 11:00 am
As the fin folds the width of the transistor, so what used to be a free choice of width in the xy plane now is limited by the height of the fin in the z plane, this means all transistors are an integer multiple of the basic fin unit (ie width = 2*hypotenuse of the fin).

As can be seen by the last micrograph, this means wide transistors becoming arrays of the basic fin unit.

This might well be limiting for analog / mixed signal use of this form of transistor. There may well be other consequences I have not thought out, even for digital. It does look pretty dense on the poly and active area layers though.

*The Great Buana* says:
May 21, 2012 at 3:37 pm
How are the yields and bin-splits? Considering the small die-size, Intel should be able to produce more than enough Ivy Bridge CPUS. Instead, I read about rumours from taiwanese sources that there is not enough supply. In fact not all CPUs for Ultrabooks are 17 Watt TDP, but there are also some with 25 Watt, similar to AMDs Trinity which is still 32 nm. And there are even parts with > 35 Watt and talks that there are some allocation-shifts between Desktop and Notebook parts. Does anybody has similar information to share?

- *M-Roberts* says:
  May 21, 2012 at 4:52 pm
  From my marketing contact I hear ivy bridge CPUs come out of the fab much wider spread for electrical parameters as compared to historical planar devices. I hear "copy exactly" does not even work for CPU production at other fabs around the world which is why 17W ultrabook part was delayed. He explained that the leakage is worse (not better for the trigate). Reason is unlike on SOI….this structure has fundamentally high leakage at the bottom on the fin. Somehow related to fin being thick at the bottom. I hear this never be fixed in 22nm since process is already qualified

for production.

….semiwiki has good sources….

Plus, my designer friend looked at intel's 22nm layout and his assessment was intel would not be cost competitive in mobile. Lots of restrictive design rules to get this to work.

http://www.semiwiki.com/forum/content/1300-intel-trigate-trouble.html

- *The Great Buana* says:
  May 21, 2012 at 5:07 pm
  Sometimes I read that 22 nm and below requires Fully Depleted SOI because of leakage. But perhaps after the great success with HKMG years ahead of the competition at 45 nm Intel became too optimistic about 22 nm. With smaller die size and lower power consumption Intel was supposed to smash AMD but now they could even have some trouble. If I am not wrong then most chips are of mediocre quality, not good enough for mobile and perhaps also worse than Sandy Bridge for high performance Desktops. Even the next generation AMD Piledrivers could become competitive in the Desktop space because they are expected to be very overclockable while Ivy Bridge are not. However, Intel needs to produce chips with higher quality in order to earn the billions for their ROI.

  - *Jamo* says:
    May 21, 2012 at 9:51 pm
    When tested ivy vs sandy is same mother board and same system setup on a large number of samples ….we see little improvement in load power or idle power. We think this is consistent with tomshardware and data that says ivy does not over clock as well as sandy.

    Parts have high leakage and we also see large variation in leakage from one part to another.

    If intel can get parts to yield they should be ok for historical CPU market since die cost is not a large factor for $200 CPU and at present are not facing strong competition.

    BUT I think the real issue is intel does not have a lower power or cost effective 22nm mobile process technology.

  - *M-Roberts* says:
    May 21, 2012 at 10:50 pm
    Yes fully depleted is needed at 20nm to control leakage But for intel's trigate on bulk at the bottom of the fin the silicon thickness is too thick to fully deplete so it acts like regular CMOS but since the size is 22nm it has high leakage. So all leakage paths are not fully depleted so that is what people miss when they just assume Intel's trigate has lower leakage (or just repeat the intel spin).

    http://www.semiwiki.com/forum/content/1291-intel-selling-itself-short-trigate.html

    There are a few small fabless design houses using intels 22nm. I spoke to one of the design guys and he claims intels 22nm does in fact have higher leakage. He claims it is more of a performance technology. If you want to make mobile chips, RF or analog, foundry is better. He also confirmed variation is much worse.

*M-Roberts* says:
May 21, 2012 at 11:18 pm
Dr. James,

Again that you for this information. Anyway we can talk you into doing a part 2 with more pictures. Your article is the best the industry has to help us understand the trigate technology. Other items that would be helpful are

(1) Top down photos in GPU and CPU to illustrate intels restrictive design rules

(2) I-V of SRAM pull down transistor to measure leakage compared to foundry 28nm SRAM transistor (few pA range)

*The Great Buana* says:
May 22, 2012 at 3:29 am
Sorry, to all but could you provide more granularity? If you talk about mobile, do you refer to phones only or tablets or notebooks or all of them? Notebooks are most important for Intel, of course, including Ultrabooks. But were is the real trouble here? Are there not enough notebook-chips?

Secondly, I think performance-desktops are also likely to become an issue because Ivy Bridge shows no real improvement here while Piledriver/Vishera from AMD could come very close. The same should apply for Servers.

- *Jamo* says:
  May 22, 2012 at 10:33 am
  Each of us likly define "mobile" differently. I mean end markets that highly value portability and long battery life (smart phones, tablets, and some portion of laptops). Phone and tablet are obvious. But laptops are a more interesting

discussion.

I don't know how this all plays out and my views are just that, my personal interpretation of public data. Windows 8 will allow later this year for a "very mobile" laptop on Arm SOC for sub $499 (sub $399 if we don't not have to pay $90 to MS) which creates a very interesting platform for calendar, email, internet and MS Office 15 with > 12hr battery life.

How big is that part of the "mobile" laptop market? I have see all sorts of marking projections for 2013-14 laptop and this is what I think has Intel nervous and its marketing machine "fired up" to trash talk down.

For all Intel's trash talk, their mobile chips are MIA and this trigate appears to be a dud for power.

see comments by Raj Nair (ex-Intel?)
http://www.semiwiki.com/forum/content/1300-intel-trigate-trouble.html

The article is consistent with what we see in the end SOC chip market. This year and for Windows 8 ultra "mobile" laptops (we define by 1 to 10W SOC ), Intel will only offers 32nm atom chips that are not as good as Arm 40nm or 28nm chips for power. Intel will have been talking about this "Arm killer" low power trigate for a full 2 years before they sell the first 22nm Atom parts (note recently released 22nm Ivy Bridge are ~ 45W laptop CPUs). But the interesting thing is I don't this 22nm trigate (reverse engineered above) is much of a response and the market is now understanding this. At the system level, the upcoming ultrabook Ivy Bridge parts offer no improvement over the Sandy Bridge for power (both 17W). I was expecting 50% lower power and sub-10W Ivy Bridge parts but there is no indications these type of parts are on Intel's roadmap with Ivy Bridge or with Haswell in Q2/2013.

- *The Great Buana* says:
  May 22, 2012 at 10:49 am
  I have a different view on that. Just today digitimes released an article that OEMs are struggling to create affordable ARM-products for Windows8. And if the SOC is below 5 Watt or so it does not matter anymore how much below. And it is not so much the CPU-core but the interfaces that consume power. AMD will power down their Tablet chip to less than 5 Watt on 40 nm. They clock it down somewhat, the rest is being achieved by cutting of abundant interface bandwith for USB, SATA and all this. Today a common Southbridge already eats 7 Watts but by 2013 all will be SOC @ 3-4 Watts. Nobody will notice this anymore because the screen and anything else consumes much more and what realy counts is compatibility. The Tablet-market is overestimated anyway. 2/3s of all sales are for Apple-Fanboys and the rest is subsidised by Amazon, B&Ns and Telecom Carriers. Who else buys such toys? The good old PC faced much stronger competition in the past, from Apple to Commodore, Atari and even Acorn with ARM which was superior at this time, then the game computers but nothing could really hurt the Pc.Even know the crisis is not driven by tablets and ARM but by the financial crisis and this only matters. If there were not the countries like Italy and Spain with dramatic problems, there would be no pc crisis at all.

  I bought an Android Tablet only a few months ago for only 50 Dollars but I have no use for it.

*Dr chip designer* says:
May 22, 2012 at 1:58 pm
Tom has posted today some more power benchmarking on the 22nm process we are looking at. Comparing 3.5Ghz parts idle power is HIGHER for trigate (71W vs 66W). Transistor leakage is a key component of idle power.

http://www.tomshardware.com/reviews/core-i5-3570-low-power,3204-13.html#

Nice work Chipworks. Your TEMs give nice insight to what real world trigate deliver. Without your TEMs we would not of known the fins were sloped, variable and leaky.

Chipworks please keep up the reverse engineering and let us know if the sloped fins get fixed.

*V-FA prep* says:
May 26, 2012 at 12:48 pm
Great report!
The only problem with the correct interpretation of the structures is poor TEM images – the result of the weakness of the TEM samples preparation.
Contact me if you really want to see how it should look – with thickness of the TEM lamella below 20nm and amorphous layer less then 1nm. This are the requirements for the proper analysis of FinFETs.